

JUDCon

JBoss Users & Developers Conference

2012:India

EJB Configurations in JBoss AS7

Satish Kale & Ravindra Maurya

Agenda

- Overview of JBoss specific EJB configurations
- Brief introduction of command line interface (CLI) in AS7
- Configure EJB subsystem (using CLI)
- Override EJB configuration(s) per deployment

EJB configuration categorization

- Stateless bean configurations
- Stateful bean configurations
- Message driven bean configurations

Stateless bean configurations

- Bean instance pool
 - Stateless bean instances are pooled
 - An invocation on a stateless bean results in a request for a free bean instance, from pool
 - Different pool implementations decide how to return a free bean instance from the pool
 - Instance is returned back to the pool after the invocation is complete

strict-max-pool

- Default bean instance pool implementation available in JBoss AS7
- Has a strict upper bound (*max-pool-size*)
- Throws an exception on unavailability of a free bean instance, after reaching the *max-pool-size* and waiting for *instance-acquisition-timeout* time

Stateful bean configurations

- Bean instance cache
 - Stateful bean instances are cached
 - Stateful beans are backed by “sessions”.
 - Instances are mapped in a cache against a session id
 - Invocations on the stateful bean are forwarded to the right instance from this cache

Cache and passivation stores

- Caches are backed by passivation stores (where the instances will be passivated)
 - File based passivation store : Passivates to a local file store
 - Clustered passivation store : Infinispan backed passivation store

Message driven bean configurations

- Bean instance pool
 - Just like stateless session beans, message driven beans are backed by bean instance pools of their own
 - Resource adapter : A message driven bean is configured to a use a resource adapter. By default AS7 uses HornetQ (the high performing messaging implementation) resource adapter

Command Line Interface (CLI)

- (One of the) tools for managing your AS7 server
- Scriptable for repetitive configurations
- Most liked by admins who prefer command line based interaction for server management
- Tab completion support

CLI overview

- Launch the CLI

```
jpai@jpai:bin$ cd /home/jpai/jboss-as-7/bin
jpai@jpai:bin$ ./jboss-admin.sh --connect
[standalone@localhost:9999 /]
```

- CLI syntax – Consists of 3 parts: The *address*, an *operation name* and optional set of *parameters*

```
[/node-type=node-name (/node-type=node-name)*] : operation-name
[( [parameter-name=parameter-value (,parameter-name=parameter-value)*] )]
```

bean-instance-pool creation via CLI

- `[standalone@localhost:9999 /] /subsystem=ejb3/strict-max-bean-instance-pool=new-pool:add`

Creates a new strict-max-bean-instance-pool named “*new-pool*” with default values for *max-pool-size* and *instance-acquisition-timeout*

Above operation is composed of:

address = `/subsystem=ejb3/strict-max-bean-instance-pool=new-pool`

operation name = `add`

bean-instance-pool creation via CLI

- Newly created pool configuration persisted to configuration file:

```
<subsystem xmlns="urn:jboss:domain:ejb3:1.2">
...
  <pools>
    <bean-instance-pools>
      <strict-max-pool name="slsb-strict-max-pool" max-pool-size="20" instance-
acquisition-timeout="5" instance-acquisition-timeout-unit="MINUTES"/>
      <strict-max-pool name="mdb-strict-max-pool" max-pool-size="20" instance-
acquisition-timeout="5" instance-acquisition-timeout-unit="MINUTES"/>
      <strict-max-pool name="new-pool"/>
    </bean-instance-pools>
  </pools>
...

```

Thank You..

Q & A